

# Automatic Inference of DATR Theories

Petra Barg\*

Seminar für Allgemeine Sprachwissenschaft, Universität Düsseldorf,  
Universitätsstraße 1, D-40225 Düsseldorf, Germany

**Summary:** An approach for the automatic acquisition of linguistic knowledge from unstructured data is presented. The acquired knowledge is represented in the lexical knowledge representation language DATR. A set of transformation rules that establish inheritance relationships and a default-inference algorithm make up the basis components of the system. Since the overall approach is not restricted to a special domain, the heuristic inference strategy uses criteria to evaluate the quality of a DATR theory, where different domains may require different criteria. The system is applied to the linguistic learning task of German noun inflection.

## 1. Introduction

The following paper presents an approach for automatic acquisition of linguistic knowledge from observations within a given domain; it thus addresses a topic from the field of machine learning (cf. Michalski (1986)), or more precisely, from the field of machine learning of natural language, as it is sometimes called (cf. Powers and Reeker (1991)).

The last decade has seen a growing interest in the application of machine learning to different kinds of linguistic domains (cf. Powers and Reeker (1991)) and has been motivated by different objectives, such as modelling cognitive processes or overcoming the knowledge-acquisition bottleneck in natural-language processing systems. The principal motivation for our approach is theoretical. The automatically induced analyses can be compared to existing linguistic descriptions and thus can confirm proposed analyses or provide alternative representations. On the other hand, the approach can be used by descriptive linguists as a tool to obtain a rough structuring of an entirely new domain (e.g. a language that has not yet been investigated).

Theoretical approaches and implemented systems cover subjects from many different linguistic areas and use different kinds of learning strategies. Some are specially designed for a particular task, while others are more general and can thus be applied to other tasks as well. Whereas the present paper pursues a general approach that is not restricted to a specific linguistic domain, the system is crucially determined by the properties of the chosen representation language.

---

\* The presented work was partly supported by the Deutsche Forschungsgemeinschaft (DFG). For helpful comments and discussions on the topic we would like to thank Gerald Gazdar, Dafydd Gibbon, James Kilbury, Ingrid Renz, and Markus Walther.

In designing a learning system the choice of a language for representing the acquired knowledge is crucial for the quality of the output of the learning task. One requirement for a linguistic representation formalism is that the information can be structured in a way that captures generalizations over linguistic objects in order to minimize redundancy. Since many linguistic generalizations have exceptions that can only be treated adequately if the representation formalism includes some device for handling default information, we have chosen the language DATR (cf. Evans and Gazdar (1989), (1990)), which allows regularities and sub-/irregularities to be expressed in a uniform way.

We briefly summarize the main features of DATR here but presuppose a basic familiarity with the language as described in (Evans and Gazdar (1989)). DATR is a declarative formalism for the definition of inheritance networks. It includes orthogonal multiple inheritance and a default-inheritance mechanism. A network description in DATR is called a *theory* and describes a set of objects (*nodes*). The properties of an object are defined by *path-definition* pairs, where a path consists of an ordered sequence of atoms (enclosed in angle brackets). The definition can be either the directly stated *value* (atomic value or sequence of atomic values) of the property, an *inheritance descriptor* that states where the value of that property can be inherited from, or a *sequence* of inheritance descriptors. An inheritance descriptor can refer to another node, path or node-path pair of the theory. The triple consisting of a node, a path, and a definition is called a *definitional sentence*.

The simple DATR theory in Fig. 1 encodes information about English verb morphology. It contains the three node definitions VERB, Love, and Come, which contain three, two, and four definitional sentences, respectively. The node definition VERB encodes the information that all past tense forms of a verb are like the root plus *\_ed*, and all present tense forms are like the root, with the exception of the form for third singular. As a regular verb Love inherits all information except the morphological root from the node VERB. In contrast Come deviates from the regular verbs in its past tense forms, which are therefore specified in the node definition. All other information can be inherited from VERB.

**Fig. 1** *a simple DATR theory*

```

VERB: <mor past> == ("<mor root>" _ed)
      <mor pres tense> == "<mor root>"
      <mor pres tense sing three> == ("<mor root>" _s).
Love: <> == VERB
      <mor root> == love.
Come: <> == VERB
      <mor root> == come
      <mor past> == came
      <mor past participle> == <mor root>.

```

The information expressed in a DATR theory is accessed by *queries* concerning objects and their properties. A query consists of a node-path pair

and returns an atomic value (or a sequence of atomic values) or fails. Seven inference rules and a default mechanism are given to deterministically evaluate the queries. The query `Love:<mor pres tense sing two>` evaluates to `love` for the theory in Fig. 1. A query together with its returned value is called an *extensional sentence*.

## 2. Inference of DATR theories

Many learning systems use the same formal language to represent the input data and the acquired knowledge. Extensional sentences (which constitute the output of the conventional inference in DATR) form a natural sub-language of DATR which is suitable to represent the input data. Since extensional sentences all have atomic values and thus are not related to each other, they can be taken as representing independent and unstructured facts about a given linguistic domain. The learning task then consists in forming a DATR theory which accounts for the observed facts through adequate structuring.<sup>1</sup>

For an acquired DATR theory to be regarded as adequately characterizing a given set of observations it has to meet at least the following criteria (in addition to the general syntactic wellformedness conditions that hold for every DATR theory):

- consistency with respect to the input data
- completeness with respect to the input data
- structuring of the observed data by inheritance relationships
- structuring of the observed data by generalizing them

The first two of these criteria constitute minimal, formal requirements that can be verified easily. A DATR theory is consistent with respect to a given set of extensional sentences if, for every query that constitutes the left-hand side of one of the extensional sentences, the returned value is that of the extensional sentence. If this holds for all left-hand sides of the extensional sentences the theory is also complete with respect to the input data.

The last two criteria rely more on intuitions and cannot be checked so easily. The inferred DATR theory should structure the observed data so that it reveals relationships that exist between the extensional sentences. A DATR theory expresses such relationships by the use of inheritance descriptors.

The generalization of the observed data is twofold. First of all, a set of specific facts should be generalized, whenever this is possible, to a single more general assumption that covers all of the specific facts. In DATR such generalizations are captured by defaults expressed in sentences that

---

<sup>1</sup>Light (1994) addresses a related topic, the insertion of a new object (described with extensional DATR sentences) into an existing DATR theory. In contrast to our approach the assumption of a structured initial theory is made.

cover more than one property of an object (as opposed to the input data, where each sentence is supposed to represent a single observed property). For example, the sentence `VERB:<mor past> == ("<mor root>" _ed)` of the theory in Fig. 1 covers all past tense forms of a verb. In addition to this process of generalization which is used in many machine-learning systems (e.g. Mitchell (1982), Michalski (1983)), acquired DATR theories should identify information that several objects have in common. This information should be abstracted and stored in more general objects from which the others inherit. Such generalized objects further structure the domain because hierarchies evolve where objects are grouped into classes.

## 2.1 Acquisition of inheritance relationships

The observed data constitute a trivial DATR theory which forms the initial hypothesis  $H_0$  of the learning task. This DATR theory is complete and consistent with respect to the input but does not meet the other two criteria. This section addresses the question of how a given DATR theory can be transformed into another theory that contains more inheritance descriptors or changes the latter in order to structure the domain.

The knowledge of how a given DATR theory can be transformed into a new one with different inheritance descriptors is defined by rewrite rules of the following format:

**Fig. 2** *form of a transformation rule*

$$s_i \rightarrow s_i' / c_1, \dots, c_n$$

where  $s_i$  is the input sentence and  $s_i'$  is the transformed sentence. Since inheritance descriptors are stated as right-hand sides (RHS) or parts of RHSs of sentences, the transformation rules operate on RHSs of DATR sentences. Thus,  $s_i'$  differs from  $s_i$  in that it contains a different RHS.  $c_1, \dots, c_n$  are constraints that define under what conditions a given sentence can be transformed into another one. In order to carry out a transformation that maintains the completeness and consistency of the theory a major constraint for the application of most transformation rules to a hypothesis  $H_i$  consists in the requirement that  $H_i$  contain another sentence with the same RHS as the sentence that is to be transformed.

Corresponding to the different kinds of inheritance relationships that can be expressed in a DATR theory, there are four major groups of transformation rules: rules that return sentences with local descriptors (local paths, local nodes, local node-path pairs), rules that transform sentences into others that have a global descriptor, rules where the transformed sentence contains a descriptor that refers to a sentence with a global descriptor, and rules that create new, abstract sentences for the acquisition of a hierarchy.<sup>2</sup> In Fig. 3 the rule for creating local node descriptors is formulated. Here,  $H_i$  is the given DATR theory,  $V_a$  is the set of atomic values in  $H_i$ , and  $N$  is the set of nodes in  $H_i$ . The rule transforms a sentence  $s$  with atomic value into one

---

<sup>2</sup>Barg (1995) gives a full account of all transformation rules.

with a node descriptor  $v'$ , if the theory contains another sentence  $s_i$  that belongs to node  $v'$  and has the same path and value as  $s$ .

**Fig. 3** *rule for local node inheritance*

$$\begin{aligned} s : (n, p, v) \rightarrow s' : (n, p, v') \quad / \quad & v \in V_a, \\ & v' \in N, \\ & s_i : (v', p, v) \in H_i, \\ & s_i \neq s \end{aligned}$$

By means of transformation rules all the different kinds of inheritance descriptors can be obtained with the exception of evaluable paths. Evaluable paths capture dependencies between properties with different values and therefore cannot be acquired by transformation rules that crucially depend on the existence of sentences which have the same RHSs. Therefore, they have here been excluded from the learning task.

## 2.2 Acquisition of default information

While inheritance relationships are represented with the RHSs of sentences, default information is basically expressed through paths of the left-hand sides (LHSs), namely by paths that cover more than one fact. Since transformation rules leave the LHSs of sentences unchanged, an additional device is necessary that operates on LHSs of sentences. For this purpose a default-inference algorithm (DIA) was developed that reduces any given DATR theory that does not (yet) contain default information, where "reduction" means shortening the paths of sentences (by cutting off a path suffix) or deletion of whole sentences. Since extensive generalization is normally a desirable property, the resulting theory must be (and indeed is) maximally reduced.

In order to acquire a DATR default theory that remains consistent with respect to the input data the DIA has to check that a reduction of a sentence does not lead to any conflicts with the remaining sentences of the theory. Conflicts can only arise between sentences which have the same node and path, because in all other cases the longest matching path can be determined. Therefore, if a given sentence is to be shortened, it has to be checked whether the theory already contains another sentence with the same node and shortened path. If it does, and if the other sentence has a different RHS, the first sentence cannot be shortened and must remain in the resulting theory. If the other sentence has the same RHS, the first sentence can be removed from the theory altogether. If the theory does not contain the shortened sentence, the shortening is a legitimate operation since no conflicts can arise.

The following additional restrictions must be imposed to guarantee a theory that is complete and consistent with respect to the input data. First of all, the sentences of a node have to be considered in descending order according

to the length of their paths. This guarantees that for every sentence, the sentences it can conflict with are still contained in the theory and are not shortened or removed. For similar reasons, sentences can only be shortened by one element (the last) at a time. In the case of path references or node-path pairs, some additional tests are carried out since potential conflicts arise from DATR's mechanism governing the inheritance of path extensions.

## 2.3 Inference strategy

The inference strategy determines how a result hypothesis  $H_R$  is acquired from an initial hypothesis  $H_0$ . It relies on the notion of a permissible derivation which arises through applications of transformation rules and DIA. A permissible derivation of  $H_0$  results from any sequence of transformation rules followed by the DIA. For reasons of consistency it is not possible to apply transformation rules after the DIA or to apply the DIA several times.

Many different theories can be derived from  $H_0$ , but only some of them can be regarded as good DATR theories with respect to the input data.<sup>3</sup> In order to acquire a good theory the space of permissible derivations has to be searched. Since an exhaustive search leads to a combinatorial explosion for every non-trivial problem, a heuristic search is used as in many other systems. We employ a forward pruning strategy that works as follows: First of all, by further restricting the transformation rules and DIA, not all of the possible successor hypotheses are generated for a given hypothesis. Most importantly, the rules for building hierarchies are restricted in order to gain sensible classes. Here the notion of similarity of objects (i.e. the number of sentences that two objects have in common) plays a crucial role as in clustering approaches (cf. Stepp and Michalski (1986), Lebowitz (1987)).

Of the generated successor hypotheses only the few most promising ones are further expanded, while all others are discarded from the search. To decide which hypotheses are promising, criteria are needed to evaluate DATR theories. Since only monotonic DATR theories can be further transformed, these criteria have to be formulated for such theories. On the other hand, only default theories are considered as possible solutions, since the representation of default information constitutes a major demand on an appropriate theory. Therefore, the default theories resulting from the most promising monotonic theories are the candidates for the result hypothesis. Again, criteria are needed in order to select the best of these candidates. The search terminates when no more transformation rules can be applied.<sup>4</sup>

Each kind of criteria forms a complex that is composed of various different single criteria that are ordered according to priority. As the inference

---

<sup>3</sup>The question of what constitutes a good DATR theory is addressed later. Assume for the moment that it has been defined and that two DATR theories can be compared with each other with respect to quality.

<sup>4</sup>This presupposes that the search space is finite, which is guaranteed by further restricting the transformation rules (more precisely the rules for creating abstract sentences).

strategy is not restricted to any specific domain, different learning tasks usually require different evaluation criteria or different orderings. Among the criteria that were found to be most useful are the following:

- size of a DATR theory, measured by the absolute or average number of sentences per object (useful only for default theories)
- homogeneity of RHSs, measured by the number of different RHSs
- complexity of RHSs (length of paths and sequences)
- capturing of particular relationships such as
  - relationships between objects (relative number of node references)
  - relationships within objects (relative number of path references)

### 3. Inference of German noun inflection

An implementation of the approach has been applied to a number of different learning tasks, including the acquisition of German noun inflection (cf. Wurzel (1970)). The input for these tasks can be drawn from sample evaluations of a corresponding DATR theory that is included in the DATR papers (cf. Evans and Gazdar (1990)). It consists of sentences whose paths contain attributes for case and number and whose values are the inflected word forms (here, abstract morphemes) associated with them, as illustrated in Fig. 4. In addition, information about the root form and the gender are included.

**Fig. 4** *input sentence for German noun inflection*

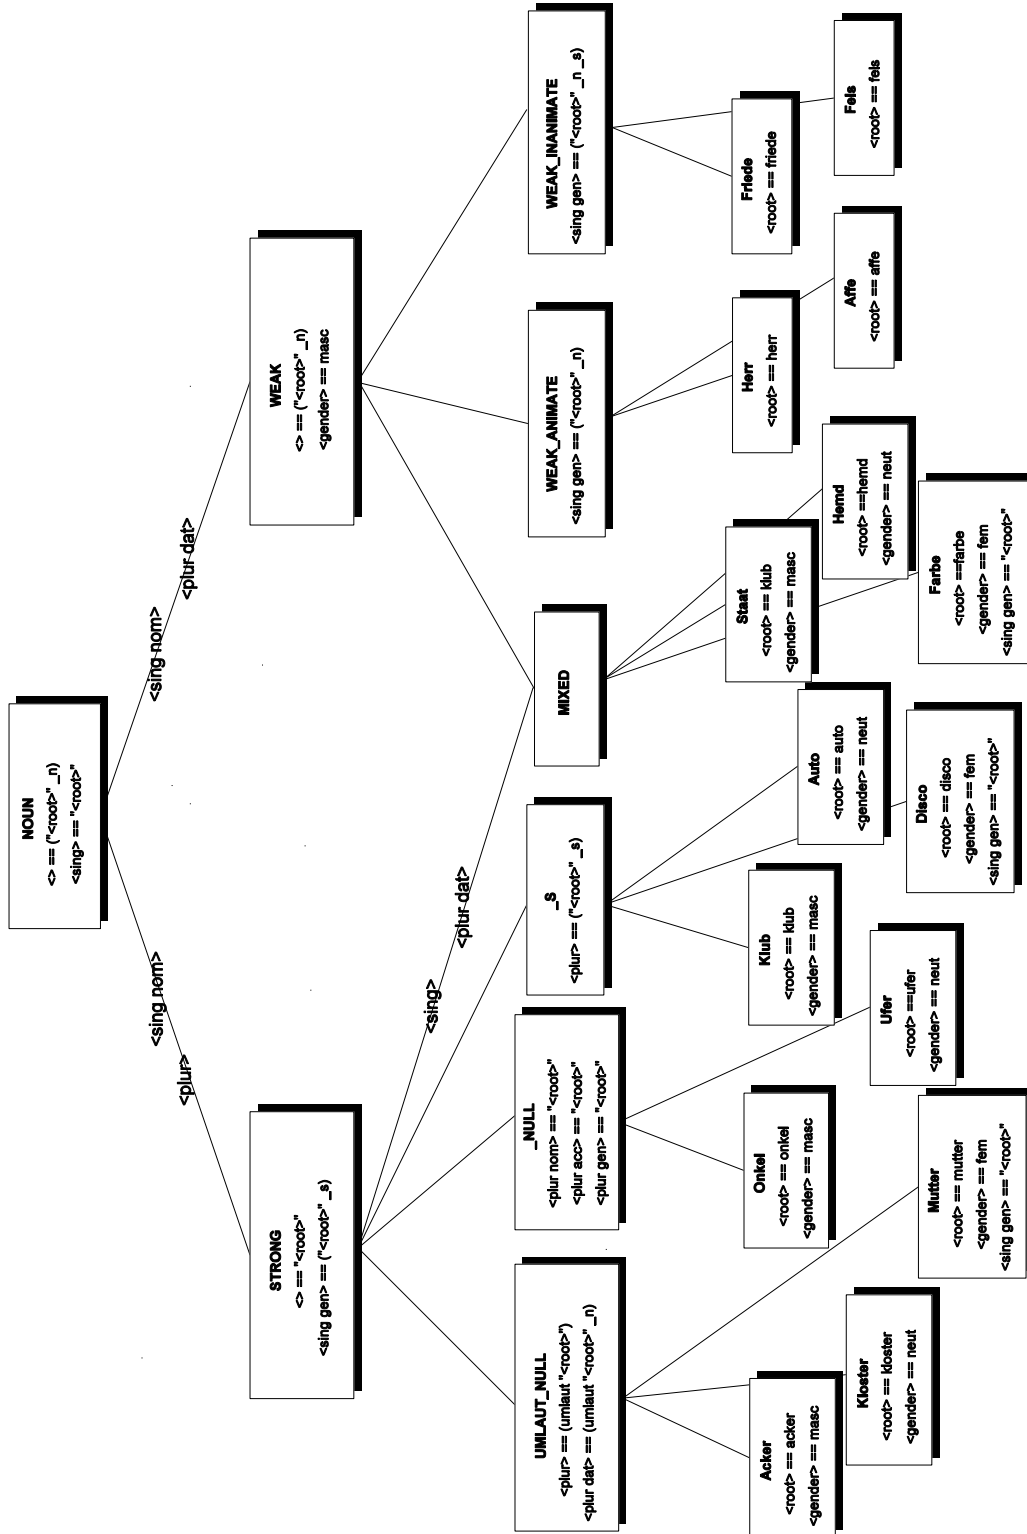
Fels: <plur nom> = (fels \_n).

For the learning task observations about nouns of various inflectional classes are given: *Fels* 'rock', *Friede* 'peace', *Herr* 'gentleman' and *Affe* 'monkey' are weak nouns, *Staat* 'state', *Hemd* 'shirt' and *Farbe* 'colour' are mixed, and *Acker* 'field', *Kloster* 'convent', *Mutter* 'mother', *Onkel* 'uncle', *Ufer* 'shore', *Klub* 'club', *Auto* 'car', and *Disco* 'disco' are strong.

The criteria for selecting the most promising theories during search were the number of different references, followed by the complexity of inheritance descriptors and number of levels in the hierarchy. The criteria for determining the best hypotheses were the number of sentences with a node-path pair on the RHS, the relative number of sentences with no node reference, and the average size of objects. All of the mentioned criteria were to be minimized. The acquired DATR theory is depicted graphically in Fig. 5. Here, the automatically generated abstract node names are replaced (manually) by more linguistically motivated names. Edges that are not annotated correspond to inheritance via the empty path.

The inferred hierarchy in Fig. 5 structures the domain of German noun

Fig. 5 *acquired hierarchy for German noun inflection*





inflection in a linguistically plausible way. According to similarity nouns are grouped into six major classes, from which they inherit most of their information. The first three of them (UMLAUT\_NULL, NULL, \_S) correspond to strong classes that have in common the formation of singular forms but differ in their plural forms, which are therefore stated explicitly. The last two classes (WEAK\_ANIMATE, WEAK\_INANIMATE) represent weak noun classes that differ only in the formation of their forms for genitive singular. The commonalities of strong nouns on the one hand and weak nouns on the other hand are further abstracted from these classes and specified in the two more general node definitions STRONG and WEAK respectively. As an interesting fact, the class of mixed nouns (MIXED) has been identified, whose members behave like strong nouns with respect to the formation of their singular forms and like weak nouns in the formation of their plural forms. These facts are captured by inheriting information from the classes STRONG and WEAK. Finally, the top node NOUN of the hierarchy represents information that is typical for German nouns in general.

## 4. Conclusion

This paper has presented an approach to the acquisition of linguistic knowledge from unstructured data. The approach is general in the sense that it is not restricted to a specific linguistic domain. This has been achieved by choosing the general representation language DATR for the representation of the acquired knowledge and by postulating a learning strategy that is tailor-made for this formalism. A similar approach could be conceived for other knowledge-representation formalisms (e.g. KL-ONE, cf. Brachman and Schmolze (1985)) which are more familiar within the artificial-intelligence paradigm.

The system was applied to a learning task involving German noun inflection. The results are sensible in that nouns are grouped into classes according to their inflectional behavior in such a way that generalizations are captured. The acquired theories are restricted in that they do not make use of evaluable paths; thus, although they are clearly non-trivial, the theories constitute a proper sublanguage of DATR. In the future, further applications of the system within different domains must be made in order to get a more detailed view of its possibilities. This pertains especially to the criteria for guiding the search and selecting best hypotheses.

## References:

- BARG, P. (1995): Automatischer Erwerb von linguistischem Wissen: ein Ansatz zur Inferenz von DATR-Theorien. Dissertation, Heinrich-Heine-Universität Düsseldorf.
- BRACHMAN, R.J., and SCHMOLZE, J.G. (1985): An Overview of the KL-ONE Knowledge Representation System. *Cognitive Science*, 9, 171-216.

- EVANS, R., and GAZDAR, G. (1989): Inference in DATR. *Proc. of the 4th Conference of the European Chapter of the Association for Computational Linguistics*, 66-71.
- EVANS, R., and GAZDAR, G. (eds.) (1990): *The DATR Papers: February 1990 (= Cognitive Science Research Paper 139)*. School of Cognitive and Computing Sciences, University of Sussex, Brighton, England.
- LEBOWITZ, M. (1987): Experiments with Incremental Concept Formation: UNIMEM. *Machine Learning*, 2, 103-138.
- LIGHT, M. (1994): Classification in Feature-based Default Inheritance Hierarchies. In H. Trost (ed.): *KONVENS '94: Verarbeitung natürlicher Sprache*. Österreichische Gesellschaft für Artificial Intelligence, Wien, 220-229.
- MICHALSKI, R. (1983): A Theory and Methodology of Inductive Learning. *Artificial Intelligence*, vol. 20, 2, 111-161.
- MICHALSKI, R. (1986): Understanding the nature of learning: Issues and research directions. In R.S. Michalski, J.G. Carbonell, and T.M. Mitchell (eds.): *Machine Learning: An Artificial Intelligence Approach*. Los Altos, Morgan Kaufmann, vol. 2, 3-25.
- MITCHELL, T.M. (1982): Generalization as search. *Artificial Intelligence*, vol. 18, 203-226.
- POWERS, D. and REEKER, L. (1991): *Machine Learning of Natural Language and Ontology (Proc. AAAI Spring Symposium)*. Kaiserslautern.
- STEPP, R.E. and MICHALSKI, R.S. (1986): Conceptual Clustering: Inventing Goal-Oriented Classifications of Structured Objects. In R.S. Michalski, J.G. Carbonell, and T.M. Mitchell (eds.): *Machine Learning: An Artificial Intelligence Approach*. Los Altos, Morgan Kaufmann, vol. 2, 471-498.
- WURZEL, W. (1970): *Studien zur deutschen Lautstruktur*. Akademie-Verlag, Berlin.